

GENERATION OF MOTION BLUR

The invention relates to a method of generating motion blur in a graphics system, and to a graphics computer system. shows the temporal pre-filtering using stretched texels in accordance with an embodiment of the invention, based on displacement vector SDV and

5

Usually, images are displayed on a display screen of a display apparatus in successive frames of lines. 3D objects displayed on the display screen which move with a large speed have a large frame to frame displacement. This is in particular the case for 3D 10 games. The large displacement may lead to visual artifacts, often referred to as temporal aliasing. Temporal filtering, which adds blur to the images, alleviates these artifacts.

An expensive approach to alleviate temporal aliasing is to increase the frame rate such that the motions of the objects result in smaller frame to frame displacements. However, a high refresh rate requires an expensive display apparatus capable to display 15 images with these high refresh rates.

Another approach is temporal super-sampling wherein the images are rendered multiple times within the frame display time interval. The rendered images are averaged and then displayed. This approach requires the 3D application to send the geometry for several instances within the frame to frame interval which requires a very powerful processing.

20 A cost effective solution is to average a present image during the present frame with the previous displayed image of the preceding frame. This approach provides an approximation of motion blur only, it does not provide a satisfactory quality of the images.

US-B-6,426,755 discloses a graphics system and method for performing blur effects. In one embodiment, the system comprises a graphics processor, a sample buffer, and 25 a sample-to-pixel calculation unit. The graphics processor is configured to render a plurality of samples based on a set of received three-dimensional graphics data. The processor is also configured to generate sample tags for the samples, wherein the sample tags are indicative of whether or not the samples are to be blurred. The super-sampled sample buffer receives and stores the samples from the graphics processor. The sample-to-pixel calculation unit receives

and filters the samples from the super-sampled sample buffer to generate output pixels which form an image on a display device. The sample-to-pixel calculation units are configured to select the filter attributes used to filter the samples into output pixels based on the sample tags.

5

It is an object of the invention to add the blur during a rasterization operation with a one-dimensional filter.

10 A first aspect of the invention provides a method of generating motion blur in a graphics system as claimed in claim 1. A second aspect of the invention provides a computer graphics system as claimed in claim 13. Advantageous embodiments are defined in the dependent claims.

15 In the method of generating motion blur in a graphics system in accordance with the first aspect of the invention, geometrical information defining a shape of a graphics primitive is received, this geometrical information may be the three-dimensional graphics data referred to in US-B-6,426,755. It is also possible to use two-dimensional graphics data which is supplied by an application in a system which has less processing resources. The method uses displacement information determining a displacement vector defining a direction of motion of the graphics primitive to sample the graphics primitive in the direction 20 of the motion to obtain input samples. A one dimensional spatial filtering of the input samples provides the temporal filtering. In this manner a high quality blur is obtained without requiring complex processing and filtering.

25 A simple one dimensional filter is used without requiring redundant calculations. In contrast, the post-processing of US-B-6,426,755 has to calculate a two-dimensional filter with a per pixel varying direction and amount of filtering. The approach in accordance with the invention has the advantage that sufficient motion blur is introduced in an effective manner. It is not required to increase the frame rate, nor to increase the temporal sample rate, the quality of the images is better than obtained by the prior art averaging.

30 A further advantage is that this approach can be implemented in the well known inverse texture mapping approach as claimed in claim 6, and in the forward texture mapping approach as claimed in claim 7. The known inverse mapping approach and the forward texture mapping approach as such will be elucidated in more detail with respect to Figures 2 and 4.

In an embodiment in accordance with the invention as defined in claim 2, the footprint of the one-dimensional filter varies with the magnitude of the displacement vector and thus with the motion. This has the advantage that the amount of blur introduced is correlated with the amount of displacement of a graphics primitive. If a low amount of movement is present, only a low amount of blur is introduced and a high amount of sharpness is preserved. If a high amount of movement is present, a high amount of blur is introduced to suppress the temporal aliasing artifacts. Thus, an optimal amount of blur is provided. It is easy to vary the amount of filtering because a one-dimensional filter is required only.

In an embodiment in accordance with the invention as defined in claim 3, the displacement vector is supplied by the 2D (two-dimensional) or 3D (three-dimensional) application which, for example, is a 3D game. This has the advantage that the programmers of the 2D or 3D application have full control over the displacement vector and thus can steer the amount of blur introduced.

In an embodiment in accordance with the invention as defined in claim 4, the 2D or 3D application provides information which defines the position and the orientation of the graphics primitives during a previous frame. The method of generating motion blur in accordance with an embodiment of the invention determines the displacement vector of the graphics primitives by comparing the position and the orientation of the graphics primitives in the present frame with the position and the orientation of the graphics primitives of the previous frame. This has the advantage that the displacement vectors do not have to be calculated by the 3D application in software, but instead the geometry acceleration hardware can be used for determining the displacement vectors.

In an embodiment in accordance with the invention as defined in claim 5, the buffering of the position and the orientation of the graphics primitives during the previous frame is performed by the method of generating motion blur in accordance with the invention. This has the advantage that a standard 3D application can be used, the displacement vectors are completely determined by the method of generating motion blur in accordance with the invention.

In an embodiment in accordance with the invention as defined in claim 6, the method of generating motion blur is implemented in the well known inverse texture mapping approach.

The intensities of the pixels present in the screen space define the displayed image on the screen. Usually, the pixels are actually positioned (in a matrix display) or thought to be positioned (in a CRT) in an orthogonal matrix indicated by an orthogonal x and

y coordinate system. In the embodiment in accordance with the invention as defined in claim 6, the x and y coordinate system is rotated such that the screen displacement vector in the screen space occurs in the direction of the x-axis. Therefore, the sampling is performed in the screen space in the direction of the screen displacement vector. The graphics primitive in the screen space is the real world graphics primitive mapped (also referred to as projected) to the rotated screen space. Usually, the graphics primitive is a polygon. The screen displacement vector is the displacement vector of the eye space graphics primitive mapped to the screen space. The eye space graphics primitive is also referred to as the real world graphics primitive, which does not indicate that a physical object is meant, also synthetic objects are covered. The sampling provides coordinates of the resampled pixels which are used as input samples for the inverse texture mapping, instead of the coordinates of the pixels in the non-rotated coordinate system.

Then, the well known inverse texture mapping is applied. A blurring-filter which has a footprint in the rotated coordinate system, is allocated to the pixels. The pixels within the footprint will be filtered in accordance with the blurring-filter amplitude characteristics. The footprint in the screen space is mapped to the texture space and called the mapped footprint. Also the polygon in the screen space is mapped to the texture space and called the mapped polygon. The texture space comprises the textures which should be displayed on the surface of the polygon. These textures are defined by texel intensities stored in a texture memory. Thus, the textures are appearance information which define an appearance of the graphics primitive by defining texel intensities in a texture space.

The texels both falling within the mapped footprint and within the mapped polygon are determined, the mapped blurring-filter is used to weight the texel intensities of these texels to obtain the intensities of the pixels in the rotated coordinate system (thus, the intensities of the resampled pixels instead of the intensities of the pixels in the well known inverse texture mapping wherein the coordinate system is not rotated).

The one-dimensional filtering averages the intensities of the pixels in the rotated coordinate system to obtain averaged intensities. A resampler resamples the averaged pixel intensities of the resampled pixels to obtain the intensities of the pixels in the original non-rotated coordinate system from the averaged intensities.

In an embodiment in accordance with the invention as defined in claim 7, the method of generating motion blur is implemented in the forward texture mapping approach.

In the texture space the texel intensities of the graphics primitive in the texture space are resampled in the direction of a texture displacement vector to obtain resampled

texels (RT_i). The texel displacement vector is the real world displacement vector mapped to the texel space. The texel intensities, which are stored in a texture memory, are interpolated to obtain the intensities of the resampled texels. The one-dimensional spatial filtering averages the intensities of the resampled texels in accordance with a weighting function to obtain filtered texels. The filtered texels of the graphics primitive are mapped to the screen space to obtain mapped texels. The intensity contributions of a mapped texel to all the pixels of which a corresponding pre-filter footprint of a pre-filter covers the mapped texel is determined. The contribution of a mapped texel to a particular pixel depends on the characteristic of the pre-filter. For each pixel, the intensity contributions of the mapped texels are summed to obtain the intensity of each one of the pixels.

Thus, said in other words, the coordinates of texels within the polygon in texture space are mapped to the screen space, and a contribution from a mapped texel to all the pixels of which the corresponding pre-filter footprint covers this texel is determined in accordance with the filter characteristic for this texel, and finally all the contribution of the texels are summed for each pixel to obtain the pixel intensity.

In an embodiment in accordance with the invention as defined in claim 8, the displacement vector of the graphics primitive is determined as an average of the displacement vectors of vertices of the graphics primitive. This has the advantage that only a single displacement vector for each polygon is required, which displacement vector can be determined in an easy manner. It suffices if the directions of the displacement vectors of the vertices is averaged. The magnitude of the displacement vector may be interpolated over the polygon.

In an embodiment in accordance with the invention as defined in claim 9, the intensities of the resampled pixels are distributed, in the screen space, in a direction of the displacement vector in the screen space over a distance determined by a magnitude of the displacement vector to obtain distributed intensities. The overlapping distributed intensities of different pixels are averaged to obtain a piece-wise constant signal which is the averaged intensity in screen space. This has the advantage that a shutter behavior of a camera is resembled, thus providing a very acceptable motion blur.

In an embodiment in accordance with the invention as defined in claim 10, the intensities of the resampled texels are distributed, in the texture space, in a direction of the displacement vector in the texture space over a distance determined by a magnitude of the displacement vector to obtain distributed intensities. The overlapping distributed intensities of different resampled texels are averaged to obtain a piece-wise constant signal which is the

averaged intensity in the texture space (also referred to as filtered texel). This has the advantage that a shutter behavior of a camera is resembled, thus providing a very acceptable motion blur.

In an embodiment in accordance with the invention as defined in claim 11, the 5 one-dimensional spatial filtering applies different weighted averaging functions during one or more frame-to-frame intervals. This has the advantage that although in each frame an efficient one-dimensional filter is performed, a higher-order temporal filtering is obtained. At the rendering of the frame, only partial intensities of the pixels are calculated which have to be stored. The pixel intensities of n successive frames have to be accumulated to obtain the 10 correct pixel intensities. In this case, n is the width of the temporal filter. The higher-order filtering provides less aliasing with a same amount of blur, or, equivalently, a reduced blur with the same amount of temporal aliasing.

In an embodiment in accordance with the invention as defined in claim 12, the 15 distance over which the resampled pixels or the resampled texels are distributed is rounded to a multiple of the distance between resampled texels. This avoids a doubling of the number of resampled texels during the accumulation of the distributed intensities of the texels.

These and other aspects of the invention are apparent from and will be elucidated with reference to the embodiments described hereinafter.

20

In the drawings:

Fig. 1 elucidates a display of a real world 3D object on a display screen,

Fig. 2 elucidates the known inverse texture mapping,

25 Fig. 3 shows a block diagram of a circuit for performing the known inverse texture mapping,

Fig. 4 elucidates the forward texture mapping,

Fig. 5 shows a block diagram of a circuit for performing the forward texture mapping,

30 Fig. 6 shows a block diagram of a circuit in accordance with an embodiment of the invention,

Fig. 7 elucidates the sampling in the direction of the displacement vector in the screen space,

Fig. 8 shows a block diagram of a circuit in accordance with an embodiment of the invention comprising the inverse texture mapping,

Fig. 9 elucidates the sampling in the direction of the displacement vector in the texture space,

Fig. 10 shows a block diagram of a circuit in accordance with an embodiment of the invention comprising forward texture mapping,

5 Fig. 11 shows an embodiment of a blurring filter with a footprint,

Fig. 12 shows the determination of a displacement vector of a polygon based on the displacement vectors of vertices of the polygon,

Fig. 13 shows the temporal pre-filtering using stretched texels in accordance with an embodiment of the invention, based on displacement vector SDV and

10 Fig. 14 shows the temporal pre-filtering using stretched texels in accordance with an embodiment of the invention, based on displacement vector TDV.

Fig. 15 shows the approximation of motion blur of a camera by using the stretched texels in accordance with an embodiment of the invention.

15

Fig. 1 elucidates a display of a real world 3D object on a display screen. A real world object WO, which may be a three-dimensional object such as the cube shown, is projected on a two-dimensional display screen DS. The three-dimensional object WO has a surface structure or texture which defines the appearance of the three-dimensional object 20 WO. In Fig. 1 the polygon A has a texture TA and the polygon B has a texture TB. The polygons A and B are with a more general term also referred to as the real world graphics primitives.

The projection of the real world object WO is obtained by defining an eye or camera position ECP with respect to the screen DS. In Fig. 1 is shown how the polygon SGP 25 corresponding to the polygon A is projected on the screen DS. The polygon SGP in the screen space SSP defined by the coordinates X and Y is also referred to as a graphics primitive instead of the graphics primitive in the screen space. Thus, with graphics primitive is indicated the polygon A in the eye space, or the polygon SGP in the screen space, or the polygon TGP in the texture space, it is clear from the context which graphics primitive is 30 meant. It is only the geometry of the polygon A which is used to determine the geometry of the polygon SGP. Usually, it suffices to know the vertices of the polygon A to determine the vertices of the polygon SGP.

The texture TA of the polygon A is not directly projected from the real world into the screen space SSP. The different textures of the real world object WO are stored in a

texture map or texture space TSP defined by the coordinates U and V. For example, Fig. 1 shows that the polygon A has a texture TA which is available in the texture space TSP in the area indicated by TA, while the polygon B has another texture TB which is available in the texture space TSP in the area indicated by TB. The polygon A is projected on the texture space TA such that a polygon TGP occurs such that when the texture present within the polygon TGP is projected on the polygon A the texture of the real world object WO is obtained or at least resembled as much as possible. A perspective transformation PPT between the texture space TSP and the screen space SSP projects the texture of the polygon TGP on the corresponding polygon SGP. This process is also referred to as texture mapping. 5 Usually, the textures are not all present in a global texture space, but every texture defines its 10 own texture space.

Fig. 2 elucidates the known inverse texture mapping. Fig. 2 shows the polygon SGP in the screen space SSP and the polygon TGP in the texture space TSP. To facilitate the elucidation, it is assumed that both the polygon SGP and the polygon TGP correspond to the 15 polygon A of the real world object WO of Fig. 1.

The intensities PIi of the pixels Pi present in the screen space SSP define the image displayed. Usually, the pixels Pi are actually positioned (in a matrix display) or thought to be positioned (in a CRT) in an orthogonal matrix of positions. In Fig. 2 only a 20 limited number of the pixels Pi is indicated by the dots. The polygon SGP is shown in the screen space SSP to indicate which pixels Pi are positioned within the polygon SGP.

The texels or texel intensities Ti in the texture space TSP are indicated by the intersections of the horizontal and vertical lines. These texels Ti which usually are stored in a memory called texture map define the texture. It is assumed that the part of the texel map or texture space TSP shown corresponds to the texture TA shown in Fig. 1. The polygon TGP is 25 shown in the texture space TSP to indicate which texels Ti are positioned within the polygon TGP, these texels Ti are indicated by small crosses.

The well known inverse texture mapping comprises the steps elucidated in the now following. A blurring-filter which has a footprint FP is shown in the screen space SSP and has to operate on the pixels Pi to perform a weighted averaging operation required to 30 obtain the blurring. This footprint FP in the screen space SSP is mapped to the texture space TSP and called the mapped footprint MFP. The polygon TGP which may be obtained by mapping the polygon SGP from the screen space SSP to the texture space TSP is also called the mapped polygon. The texture space TSP comprises the textures TA, TB (see Fig. 1) which should be displayed on the surface of the polygon SGP. As described above, these

textures TA, TB are defined by texel intensities T_i stored in a texel memory. Thus, the textures TA, TB are appearance information which define an appearance of the graphics primitive SGP by defining texel intensities T_i in a texture space TSP.

5 The texels T_i both falling within the mapped footprint MFP and within the mapped polygon TGP are determined. These texels T_i are indicated by the crosses. The mapped blurring-filter MFP is used to weight the texel intensities T_i of these texels T_i to obtain the intensities of the pixels P_i .

10 Fig. 3 shows a block diagram of a circuit for performing the known inverse texture mapping. The circuit comprises a rasterizer RSS which operates in the screen space SSP, a resampler RTS in the texture space TSP, a texture memory TM and a pixel fragment processing circuit PFO. U_t, V_t is the texture coordinate of a texel T_i with index t , X_p, Y_p is the screen coordinate of a pixel with index p , I_t is the color of the texel T_i with index t , and I_p is the filtered color of pixel P_i with index p .

15 The rasterizer RSS rasterizes the polygon SGP in the screen space SSP. For every pixel P_i traversed, its blurring filter footprint FP is mapped to the texture space TSP. The texels T_i within the mapped footprint MFP and within the mapped polygon TGP are determined and weighted according to a mapped profile of the blurring filter. The color of the pixels P_i is computed using the mapped blurring filter in the texture space TSP.

20 Thus, the rasterizer RSS receives the polygons SGP in the screen space SSP to supply the mapped blurring filter footprint MFP and the coordinates of the pixels P_i . A resampler in the texture space RTS receives the mapped blurring filter footprint MFP and information on the position of the polygon TGP to determine which texels T_i are within the mapped footprint MFP and within the polygon TGP. The intensities of the texels T_i determined in this manner are retrieved from the texture memory TM. The blurring filter 25 filters the relevant intensities of the texels T_i determined in this manner to supply the filtered color I_p of the pixel P_i .

30 The pixel fragment processing circuit PFO provides an improved processing of pixels near to borders of the polygon SGP and supplies the pixel intensities P_{II} of the pixels P_i displayed on the screen DS. The pixel fragment processing circuit PFO blends the blur if several polygons are moving.

Fig. 4 elucidates forward texture mapping. Fig. 4 shows the polygon SGP in the screen space SSP and the polygon TGP in the texture space TSP. To facilitate the elucidation, it is assumed that both the polygon SGP and the polygon TGP correspond to the polygon A of the real world object WO of Fig. 1.

The intensities P_{iI} of the pixels P_i present in the screen space SSP define the image displayed. The pixels P_i are indicated by the dots. The polygon SGP is shown in the screen space SSP to indicate which pixels P_i are positioned within the polygon SGP. The pixel actually indicated by P_i is positioned outside the polygon SGP. With each pixel P_i a footprint FP of a blur filter is associated.

The texels or texel intensities T_i in the texture space TSP are indicated by the interstices of the horizontal and vertical lines. Again, these texels T_i which usually are stored in a memory called texture map define the texture. It is assumed that the part of the texel map or texture space TSP shown corresponds to the texture TA shown in Fig. 1. The polygon TGP is shown in the texture space TSP to indicate which texels T_i are positioned within the polygon TGP.

The coordinates of the texels T_i within the polygon TGP are mapped (resampled) to the screen space SSP. In Fig. 4, this mapping (indicated by the arrow AR from the texture space TSP to the screen space SSP) of a texel T_i (indicated by a cross in the texture space) to the screen space SSP provides mapped texels MT_i (indicated by the cross in the screen space SSP, which cross may be positioned in-between pixel positions indicated by the dots) in the screen space SSP. A contribution of the mapped texel MT_i to all the pixels P_i which have a footprint FP of the blur filter which encompasses the mapped texel MT_i is determined in accordance with the filter characteristic of the blur filter. All the contributions of the mapped texels MT_i to the pixels P_i are summed to obtain the intensities P_{iI} of the pixels P_i .

In the forward texture mapping, the resampling from the colors of the texel T_i to the colors of the pixels P_i occurs in the screen space SSP, and thus is input sample driven. Compared to the inverse texture mapping, it is easier to determine which texels T_i contribute to a particular pixel P_i . Only the mapped texels MT_i which are within a footprint FP of the blurring filter for a particular pixel P_i will contribute to the intensity or color of this particular pixel P_i . Further, there is no need to transform the blurring filter from the screen space SSP to the texel space TSP.

Fig. 5 shows a block diagram of a circuit for performing the forward texture mapping. The circuit comprises a rasterizer RTS which operates in the texture space TSP, a resampler RSS in the screen space SSP, a texture memory TM and a pixel fragment processing circuit PFO. U_t, V_t is the texture coordinate of a texel T_i with index, X_p, Y_p is the screen coordinate of a pixel with index p , I_t is the color of the texel T_i with index t , and I_p is the filtered color of pixel P_i with index p .

The rasterizer RTS rasterizes the polygon TGP in the texture space TSP. For every texel T_i which is within the polygon TGP, the resampler in the screen space RSS maps the texel T_i to a mapped texel M_i in the screen space SSP. Further, the resampler RSS determines the contribution of a mapped texel M_i to all the pixels P_i of which the associated footprint FP of the blurring filter encompasses this mapped texel M_i . Finally, the resampler RSS sums the intensity contributions of all mapped texels M_i to the pixels P_i to obtain the intensities P_i of the pixels P_i .

The pixel fragment processing circuit PFO provides an improved processing of pixels near to borders of the polygon SGP and supplies the pixel intensities P_i of the pixels P_i displayed on the screen DS.

Fig. 6 shows a block diagram of a circuit in accordance with an embodiment of the invention. This motion blur generating circuit comprises a rasterizer RA, a displacement providing circuit DIG, and a one-dimensional filter ODF.

The rasterizer RA receives both geometrical information GI which defines the shape of a graphics primitive SGP or TGP and displacement information DI which determines a displacement vector defining a direction of the motion of the graphics primitive SGP or TGP. The rasterizer RA samples the graphics primitive SGP or TGP in the direction of the displacement vector to obtain samples R_i . The one-dimensional filter ODF provides a temporal pre-filtering by filtering the samples R_i to obtain averaged intensities A_i .

The rasterizer RA may operate in the screen space SSP or in the texture space TSP. If the rasterizer RA operates in the screen space SSP, the graphics primitive SGP or TGP may be the polygon SGP, and the samples R_i are based on the pixels P_i . If the rasterizer RA operates in the texture space TSP, the graphics primitive SGP or TGP may be the polygon TGP, and the samples R_i are based on the texels T_i .

The use of a rasterizer RA in the screen space SSP is elucidated with respect to Fig. 7 and with respect to its combination with the inverse texture mapping (see Fig. 8).

The use of a rasterizer RA in the texture space TSP is elucidated with respect to Fig. 9 and with respect to its combination with the forward texture mapping (see Fig. 10).

Fig. 7 elucidates the sampling in the direction of the displacement vector in the screen space. The real world object WO moves in a certain direction. This movement of the complete object WO causes the graphics primitives (the polygons A and B) to move also. The movement of the polygon A can be indicated in the screen space SSP by the displacement vector SDV of the polygon SGP. Other polygons of the real world object WO may have other displacement vectors. The intensities P_i of the pixels P_i are resampled such

that resampled pixels RPi are determined which are positioned in a rectangular grid of which one direction coincides with the direction of the displacement vector SDV. The pixels Pi are indicated by dots, the resampled pixels RPi are indicated by crosses. Only a few pixels Pi and resampled pixels RPi are shown.

5 The pixels Pi of which the intensities PIi determine the image displayed are positioned in the orthogonal coordinate space defined by the orthogonal axis x and y. The resampled pixels RPi are positioned in the orthogonal coordinate space defined by the orthogonal axis x' and y'.

10 Fig. 8 shows a block diagram of a circuit in accordance with an embodiment of the invention comprising the inverse texture mapping.

15 The sampler RSS, which is the sampler RA shown in Fig. 6 which samples in the screen space SSP, samples within a polygon SGP in the direction of the displacement vector SDV of this polygon SGP to obtain resampled pixels RPi. Therefore, the sampler RSS receives the geometry of the polygon SGP and the displacement information DI from the displacement providing circuit DIG. The displacement information DI may comprise the direction in which the displacement occurs and the amount of displacement and thus may be the displacement vector SDV. The displacement vector SDV may be supplied by the 3D application, or may be determined by the displacement providing circuit DIG from the position of the polygon A in successive frames. The resampled pixels RPi occur in an 20 equidistant orthogonal coordinate space of positions which are aligned with the displacement vector SDV. Or said differently, the coordinate system x, y in the screen space is rotated such that a rotated coordinate system x', y' is obtained of which the x' axis is aligned with the displacement vector.

25 The inverse texture mapper ITM receives the resampled pixels RPi to supply intensities RIp. The inverse texture mapper ITM operates in the same manner as the well known inverse texture mapping as elucidated with respect to Figs. 2 and 3. But, instead of the coordinates of the pixels Pi, the coordinates of the resampled pixels RPi are used. Thus, the footprint FP of the filter in the screen space is now defined in the coordinate system which is aligned with the screen displacement vector. This footprint is mapped to the texture space 30 where the texels within both this mapped footprint and within the polygon are weighted according to the mapped filter characteristics to obtain the intensity of the resampled pixel RIp to which the footprint belongs.

The one-dimensional filter ODF comprises an averager AV and a resampler RSA. The averager AV averages the intensities RIp to obtain averaged intensities ARIp. The

averaging is performed in accordance with a weighting function WF. The resampler RSA resamples the averaged intensities ARI_i to obtain the intensities PI_i of the pixels Pi.

Fig. 9 elucidates the sampling in the direction of the displacement vector in the texture space. The real world object WO moves in a certain direction. This movement of the complete object WO causes the graphics primitives (the polygons A and B) to move also. The movement of the polygon A can be indicated in the texture space TSP by the displacement vector TDV of the polygon TGP. Other polygons of the real world object WO may have other displacement vectors. The intensities of the texels Ti are resampled such that resampled texels RTi are obtained which are positioned in a matrix of which one direction coincident with the direction of the displacement vector TDV. The texels Ti are indicated by dots, the resampled texels RTi are indicated by crosses. Only a few texels Ti and resampled texels RTi are shown.

The texels Ti of which the intensities determine the texture displayed are positioned in the orthogonal coordinate space defined by the orthogonal axis U and V. The resampled texels RTi are positioned in the orthogonal coordinate space defined by the orthogonal axis U' and V'. A distance DIS between two samples (texels Ti) in the texture space is indicated by DIS.

Fig. 10 shows a block diagram of a circuit in accordance with an embodiment of the invention comprising the forward texture mapping.

The sampler RTS, which is the sampler RA shown in Fig. 6 which samples in the texture space TSP, samples within a polygon TGP in the direction of the displacement vector SDV of this polygon SGP to obtain the resampled texels RTi. Therefore, the sampler RTS receives the geometry of the polygon TGP and the displacement information DI from the displacement providing circuit DIG. The displacement information DI may comprise the direction in which the displacement occurs and the amount of displacement and thus may be the displacement vector TDV. The displacement vector TDV may be supplied by the 3D application, or may be determined by the displacement providing circuit DIG from the position of the polygon A in successive frames.

The interpolator IP interpolates the intensities of the texels Ti to obtain the intensities RI_i of the resampled texels RTi.

The one-dimensional filtering ODF comprises an averager AV which averages the intensities RI_i in accordance with a weighting function WF to obtain filtered resampled texels FTi to which is also referred as filtered texels FTi.

The mapper MSP maps the filtered texels FTi within the polygon SGP (in more general also referred to as the graphics primitive) to the screen space SSP to obtain the mapped texels MTi (see Fig. 4).

The calculator CAL determines the intensity contributions of each of the 5 mapped texels MTi to each of the pixels Pi of which a corresponding pre-filter footprint FP of a pre-filter PRF (see Fig. 11) covers one of the mapped texels MTi. The intensity contributions depend on the characteristics of the pre-filter PRF. For example, if the pre-filter has a cubic amplitude characteristic and if a mapped texel MTi is very near to a pixel Pi, the contribution of this mapped texel MTi to the intensity of the pixel Pi is relatively large. If the 10 mapped texel is at the border of the footprint FP of the prefilter which is centered at a pixel Pi, the contribution of the mapped texel MTi is relatively small. If the mapped texel MTi is not within the footprint FP of the prefilter of a particular pixel Pi, this mapped texel MTi will not contribute to the intensity of the particular pixel Pi.

The calculator CAL sums all the contribution of the different mapped texels 15 MTi to the pixels Pi to obtain the intensities PIi of the pixels Pi. The intensity PIi of a particular pixel Pi only depends on the intensities of the mapped texels MTi within the footprint FP belonging to this particular pixel Pi and the amplitude characteristic of the pre-filter. Thus for a particular pixel Pi only the contributions of the mapped texels MTi within the footprint FP belonging to this particular pixel Pi need to be summed.

20 Fig. 11 shows an embodiment of a blurring filter with a footprint. The blurring filter (also referred to as pre-filter) PRF, which in Fig. 11 filters in the screen space SSP, has a footprint FP. The footprint FP is the area of the filter PRF in the x and/or y direction in which a mapped texel MTi contributes to a pixel Pi. The filter PRF is shown for a pixel Pi at a position Xp in the screen space SSP. In the example of the filter PRF shown, the footprint 25 FP is four pixel distances wide and covers in the x-direction the positions Xp-2, Xp-1, Xp, Xp+1, Xp+2. A mapped texel MTi which is mapped at the position Xm will contribute to the pixel Pi at the position Xp with the intensity of the mapped texel MTi multiplied with the filter value CO1.

30 Fig. 12 shows the determination of a displacement vector of a polygon based on the displacement vectors of vertices of the polygon. The polygon SGP in the screen space SSP has vertices V1, V2, V3, V4 to which the displacement vectors TDV1, TDV2, TDV3, TDV4, respectively, are associated. Preferably, the displacement vector TDV for all the pixels Pi within the polygon SGP is the average of the displacement vectors TDV1, TDV2, TDV3, TDV4. Thus, the displacement vectors TDV1, TDV2, TDV3, TDV4 are vectorially

added to obtain both the direction and the amplitude (after division by the number of vertices) of the displacement vector TDV.

More complex approaches are possible, for example, if the displacement vectors TDV1, TDV2, TDV3, TDV4 are largely different, the polygon may be divided in 5 smaller polygons.

Fig. 13 shows the temporal pre-filtering using stretched pixels in accordance with an embodiment of the invention. The one-dimensional filter ODF is performed by first distributing the intensities RI_p of the resampled pixels RP_i in the direction of the displacement vector SDV. The distribution of the intensity RI_p is performed in an area 10 around the associated resampled pixel RP_i such that the local intensity RI_p is spread out over this area. The dimensions of the area are determined by the magnitude of the displacement vector SDV. This spreading out of the intensity RI_p is also referred to as stretching the pixels RP_i. As an example only, Fig. 13 shows a motion displacement which is 3.25 times the 15 distance between two adjacent resampled pixels RP_i. The pixel stretching in the x' direction (see Fig. 7) is elucidated.

In Fig. 13A, the intensities RI_p of the resampled pixels RP_i are distributed or stretched as indicated by the horizontal lines indicated by DI_i. Each dot on the x'-axis indicates the position of a resampled pixel RP_i. The lines DI_i show that the intensity RI_p of each of the resampled pixels RP_i is distributed to cover another one of resampled pixels RP_i 20 both at the left hand side and at the right hand side of each of the resampled pixels RP_i.

Fig. 13B shows the average of the overlapping distributed intensities DI_i.

Fig. 14 shows the temporal pre-filtering using stretched texels in accordance with an embodiment of the invention. The one-dimensional filter ODF is performed by first distributing the intensities RI_i of the resampled texels RT_i in the direction of the 25 displacement vector TDV. The distribution of the intensity RI_i is performed in an area around the associated resampled texel RT_i such that the local intensity RI_i is spread out over this area. The dimensions of the area are determined by the magnitude of the displacement vector TDV. This spreading out of the intensity RI_i is also referred to as stretching the resampled texels RT_i. As an example only, Fig. 14 shows a motion displacement which is 3.25 times the 30 distance between to adjacent resampled texels RT_i. The texel stretching in the U' direction (see Fig. 9) is elucidated.

In Fig. 14A, the intensities RI_i of the resampled texels RT_i are distributed or stretched as indicated by the horizontal lines indicated by TD_i. Each dot on the U'-axis indicates the position of a resampled texel RT_i. The lines TD_i show that the intensity RI_i of

each of the resampled texels RTi is distributed to cover another one of resampled texels RTi both at the left hand side and at the right hand side of each one of the resampled texels RTi.

Fig. 14B shows the average FTi of the overlapping distributed intensities TDII.

The stretched texels are overlapping if the motion displacement during the 5 frame sample interval is larger than the distance between two adjacent resampled texels RTi. The piece-wise constant signal FTi which is obtained by averaging the overlapping parts of the distributed intensities TDII is a good approximation of the time-continue integration of a camera as will be explained with respect to Fig. 15. Thus, the result of the texel stretching is a blur which resembles the blur of a traditional camera. This blur is very acceptable to a 10 viewer. If the stretched texels are not overlapping due to no or a small amount of motion, no motion blur is generated and a spatial box reconstruction is obtained.

Fig. 14 illustrates the averaging of the overlapping parts of the distributed 15 intensities DIi for a motion displacement of 3.25 times the mapped texel distances. The obtained piece-wise constant signal FTi is an approximation of an integrated signal. It is possible to view the piece-wise constant signal FTi as a box reconstruction of artificial samples that represent the averaged overlapping parts. The artificial samples depend on a varying number of overlapping stretched texels. In Fig. 14, either three or four stretched texels overlap. This can be avoided by restricting the edges of the stretched texels to the 20 resampled or mapped texel positions RTi. Thus, a motion blur factor is used which is an integer multiple of the distance between resampled texels RTi.

Fig. 15 shows the approximation of motion blur of a camera by using the stretched texels in accordance with an embodiment of the invention. Fig. 15A shows a texel stretching of eight mapped texel distances. The line indicated by tb shows the positions of the 25 resampled texels RTi in the U' direction for a particular frame. The line indicated by te shows the positions of the resampled texels RTi in the U' direction for a frame succeeding the particular frame. The distributed intensities RIi are indicated by the lines TDII. The resulting piece-wise constant intensity FTi is shown in Fig. 15B. The solid lines indicated by CA show the motion blur introduced by a camera.

To conclude, in a preferred embodiment, the invention is directed to a method 30 of generating motion blur in a 3D-graphics system. A geometrical information GI defining a shape of a graphics primitive SGP or TGP is received RSS; RTS from a 3D-application. A displacement vector SDV; TDV defining a direction of motion of the graphics primitive SGP or TGP is also received from the 3D-application or is determined from the geometrical information. The graphics primitive SGP or TGP is sampled RSS; RTS in the direction

indicated by the displacement vector SDV; TDV to obtain input samples RPi, and an one dimensional spatial filtering ODF is performed on the input samples RPi to obtain temporal pre-filtering.

It should be noted that the above-mentioned embodiments illustrate rather than 5 limit the invention, and that those skilled in the art will be able to design many alternative embodiments without departing from the scope of the appended claims. For example, in many of the embodiments above, the processing of only one polygon is elucidated. In a practical application a huge amount of polygons (or more general: graphics primitives) may have to be processed for a complete image.

10 In the claims, any reference signs placed between parenthesis shall not be construed as limiting the claim. The word "comprising" does not exclude the presence of other elements or steps than those listed in a claim. The invention can be implemented by means of hardware comprising several distinct elements, and by means of a suitably programmed computer. In the device claim enumerating several means, several of these 15 means can be embodied by one and the same item of hardware.